

DreamerPro: Reconstruction-Free Model-Based Reinforcement Learning with Prototypical Representations

Fei Deng¹, Ingook Jang², Sungjin Ahn³

¹Rutgers University, ²ETRI, ³KAIST




RUTGERS

ETRI


KAIST

Introduction

- Model-based RL for continuous and discrete control



DeepMind Control (DMC)



Atari

Figure adapted from github.com/danijar/dreamerv2

Recurrent State-Space Model (RSSM^[1])

$$p(o_{1:T} \mid a_{1:T}) = \int \prod_{t=1}^T \underbrace{p(o_t \mid s_{\leq t}, a_{<t})}_{\text{image decoder}} \underbrace{p(s_t \mid s_{<t}, a_{<t})}_{\text{forward dynamics}} ds_{1:T}$$

[1] Learning Latent Dynamics for Planning from Pixels. Danijar Hafner et al., ICML 2019.

Recurrent State-Space Model (RSSM^[1])

$$p(o_{1:T} \mid a_{1:T}) = \int \prod_{t=1}^T \underbrace{p(o_t \mid s_{\leq t}, a_{<t})}_{\text{image decoder}} \underbrace{p(s_t \mid s_{<t}, a_{<t})}_{\text{forward dynamics}} ds_{1:T}$$



struggle with visual distractions^[2-5]

[1] Learning Latent Dynamics for Planning from Pixels. Danijar Hafner et al., ICML 2019.

[2] Learning Invariant Representations for Reinforcement Learning without Reconstruction. Amy Zhang et al., ICLR 2021.

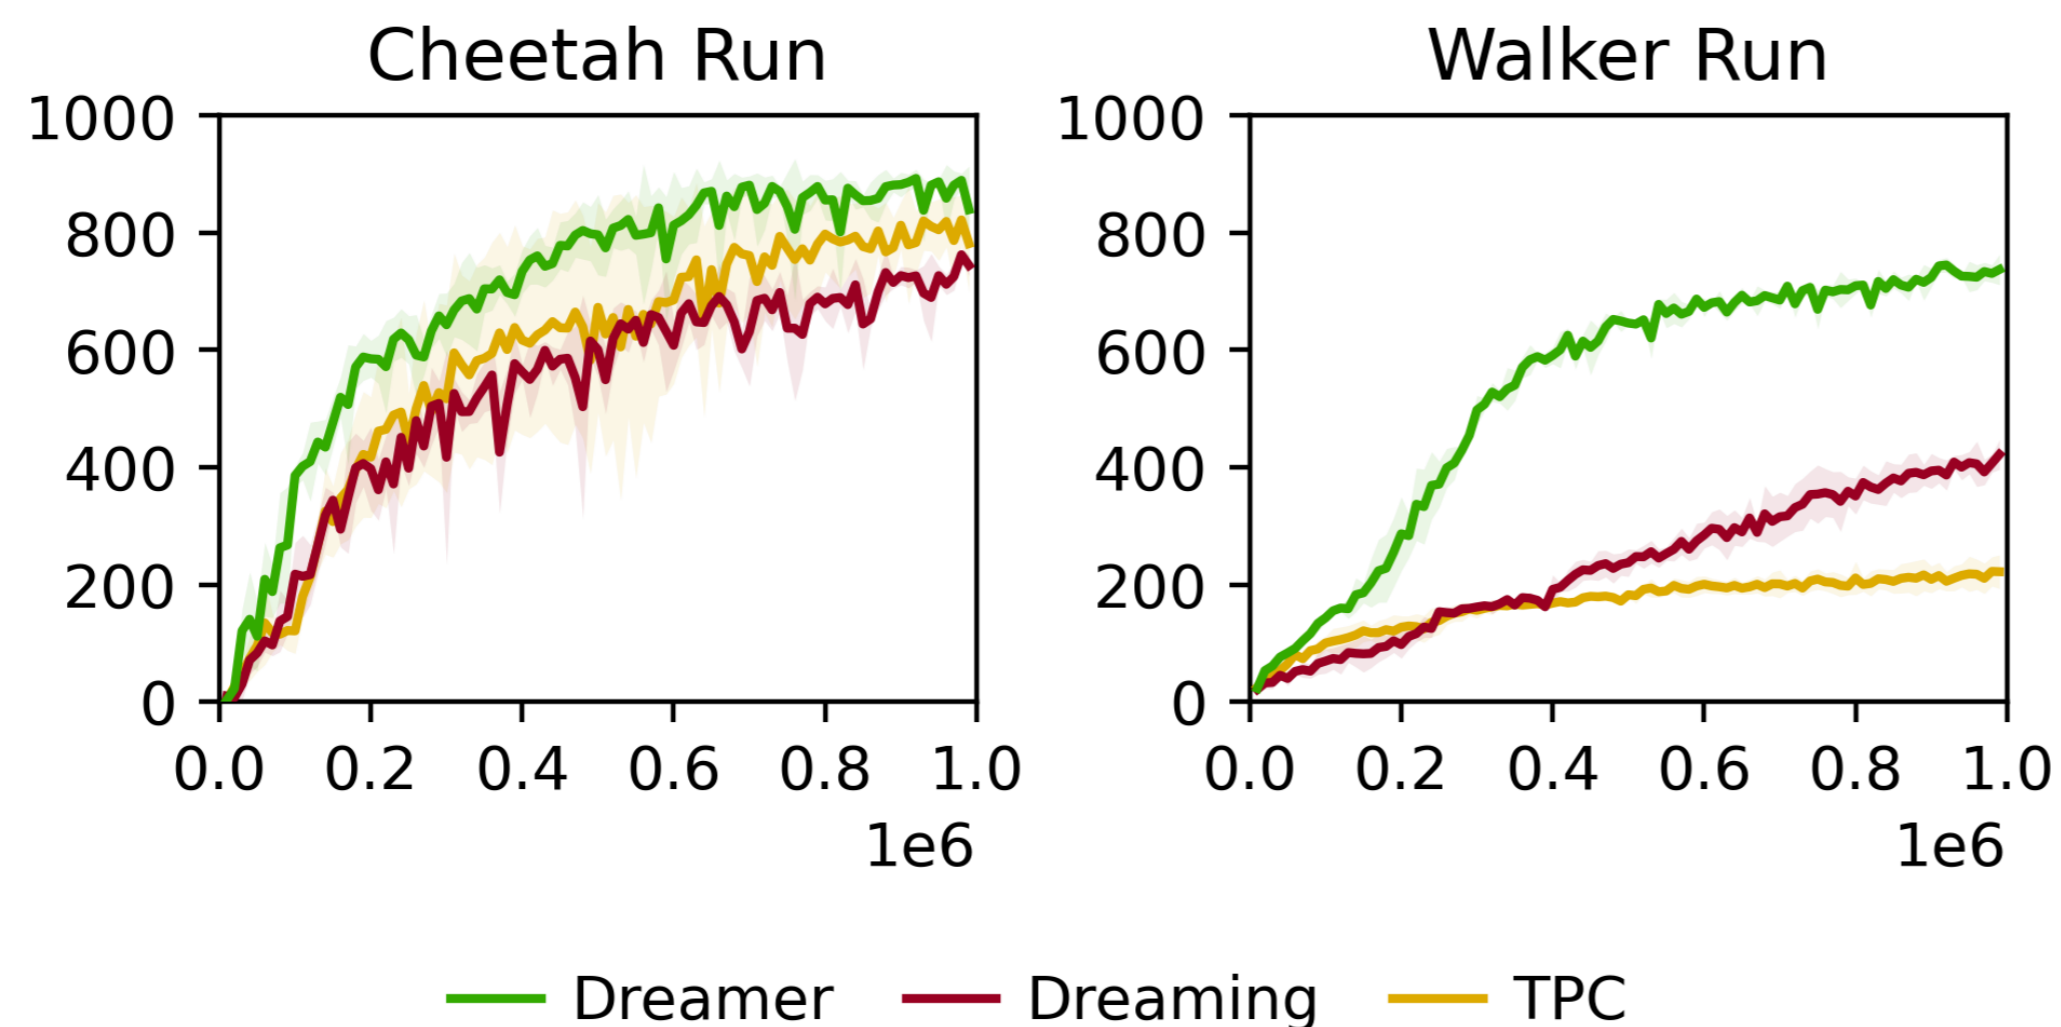
[3] Learning Task Informed Abstractions. Xiang Fu et al., ICML 2021.

[4] Temporal Predictive Coding for Model-Based Planning in Latent Space. Tung Nguyen et al., ICML 2021.

[5] Dreaming: Model-based Reinforcement Learning by Latent Imagination without Reconstruction. Masashi Okada et al., ICRA 2021.

Reconstruction-Free MBRL

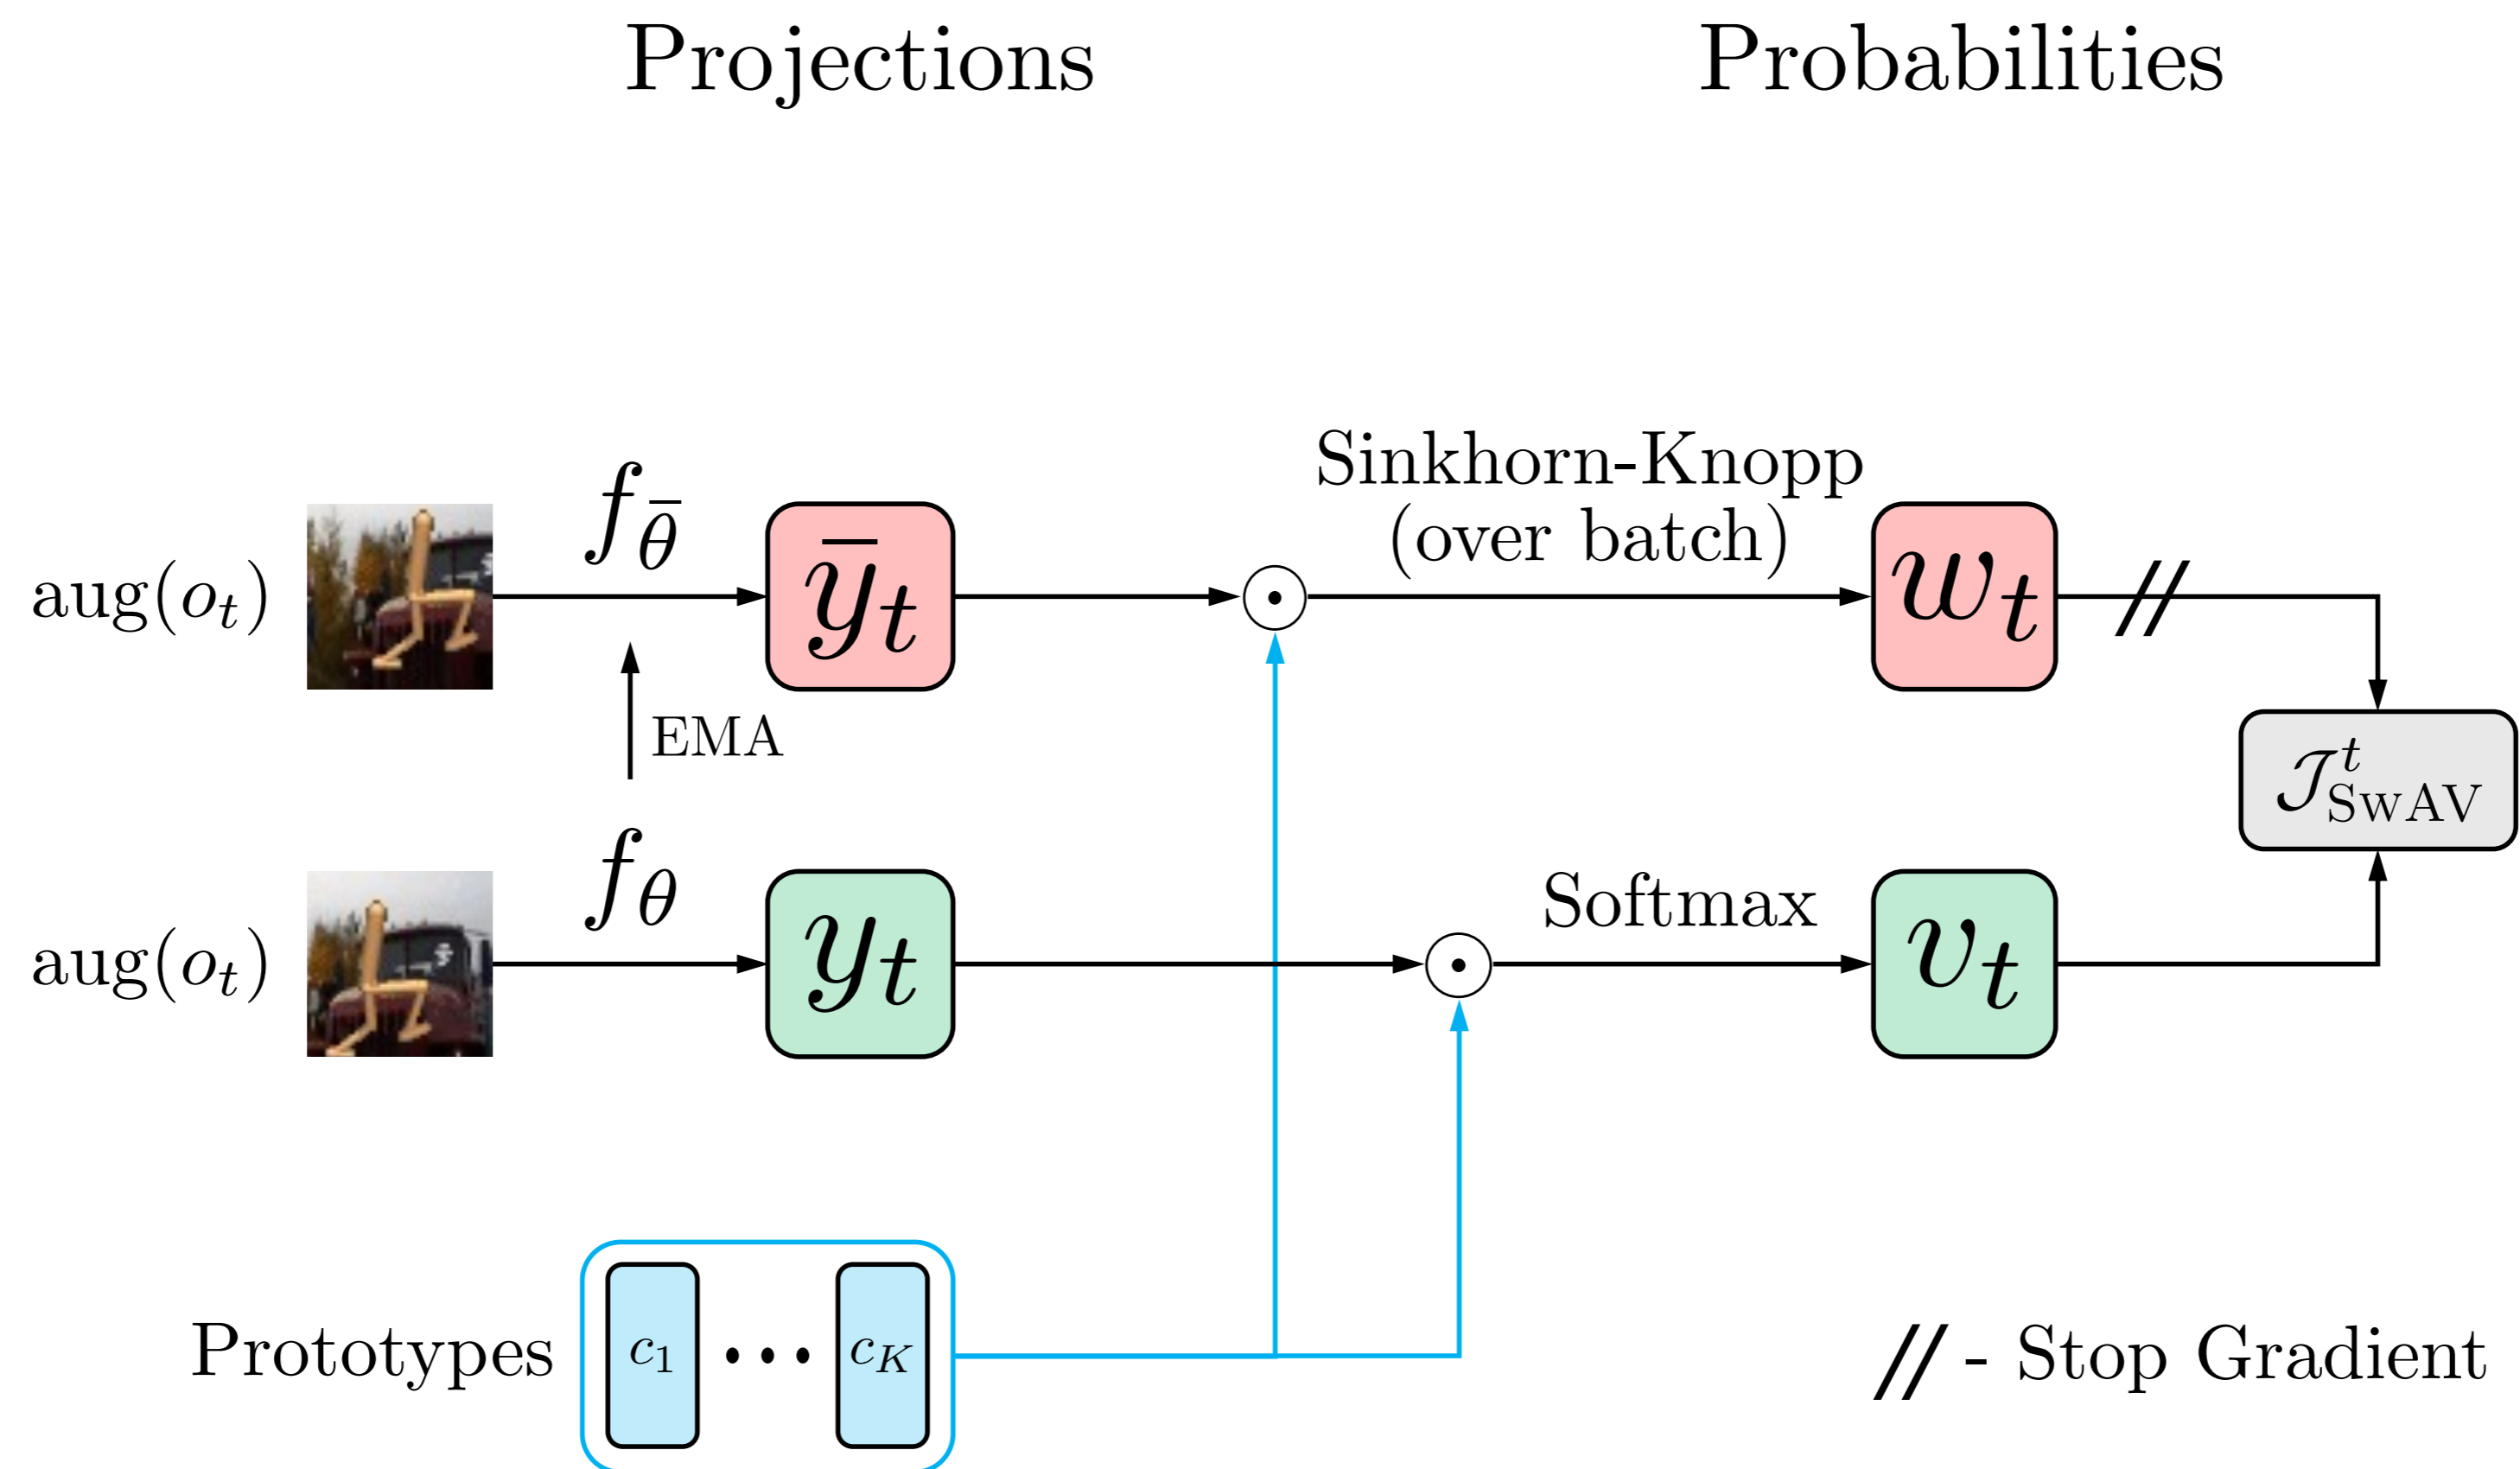
- State-of-the-art: TPC^[4], Dreaming^[5]
 - Contrastively learn the latent space
 - Underperform Dreamer in distraction-free setting



[4] Temporal Predictive Coding for Model-Based Planning in Latent Space. Tung Nguyen et al., ICML 2021.

[5] Dreaming: Model-based Reinforcement Learning by Latent Imagination without Reconstruction. Masashi Okada et al., ICRA 2021.

DreamerPro: Dreamer^[6,7] + Prototypes^[8]

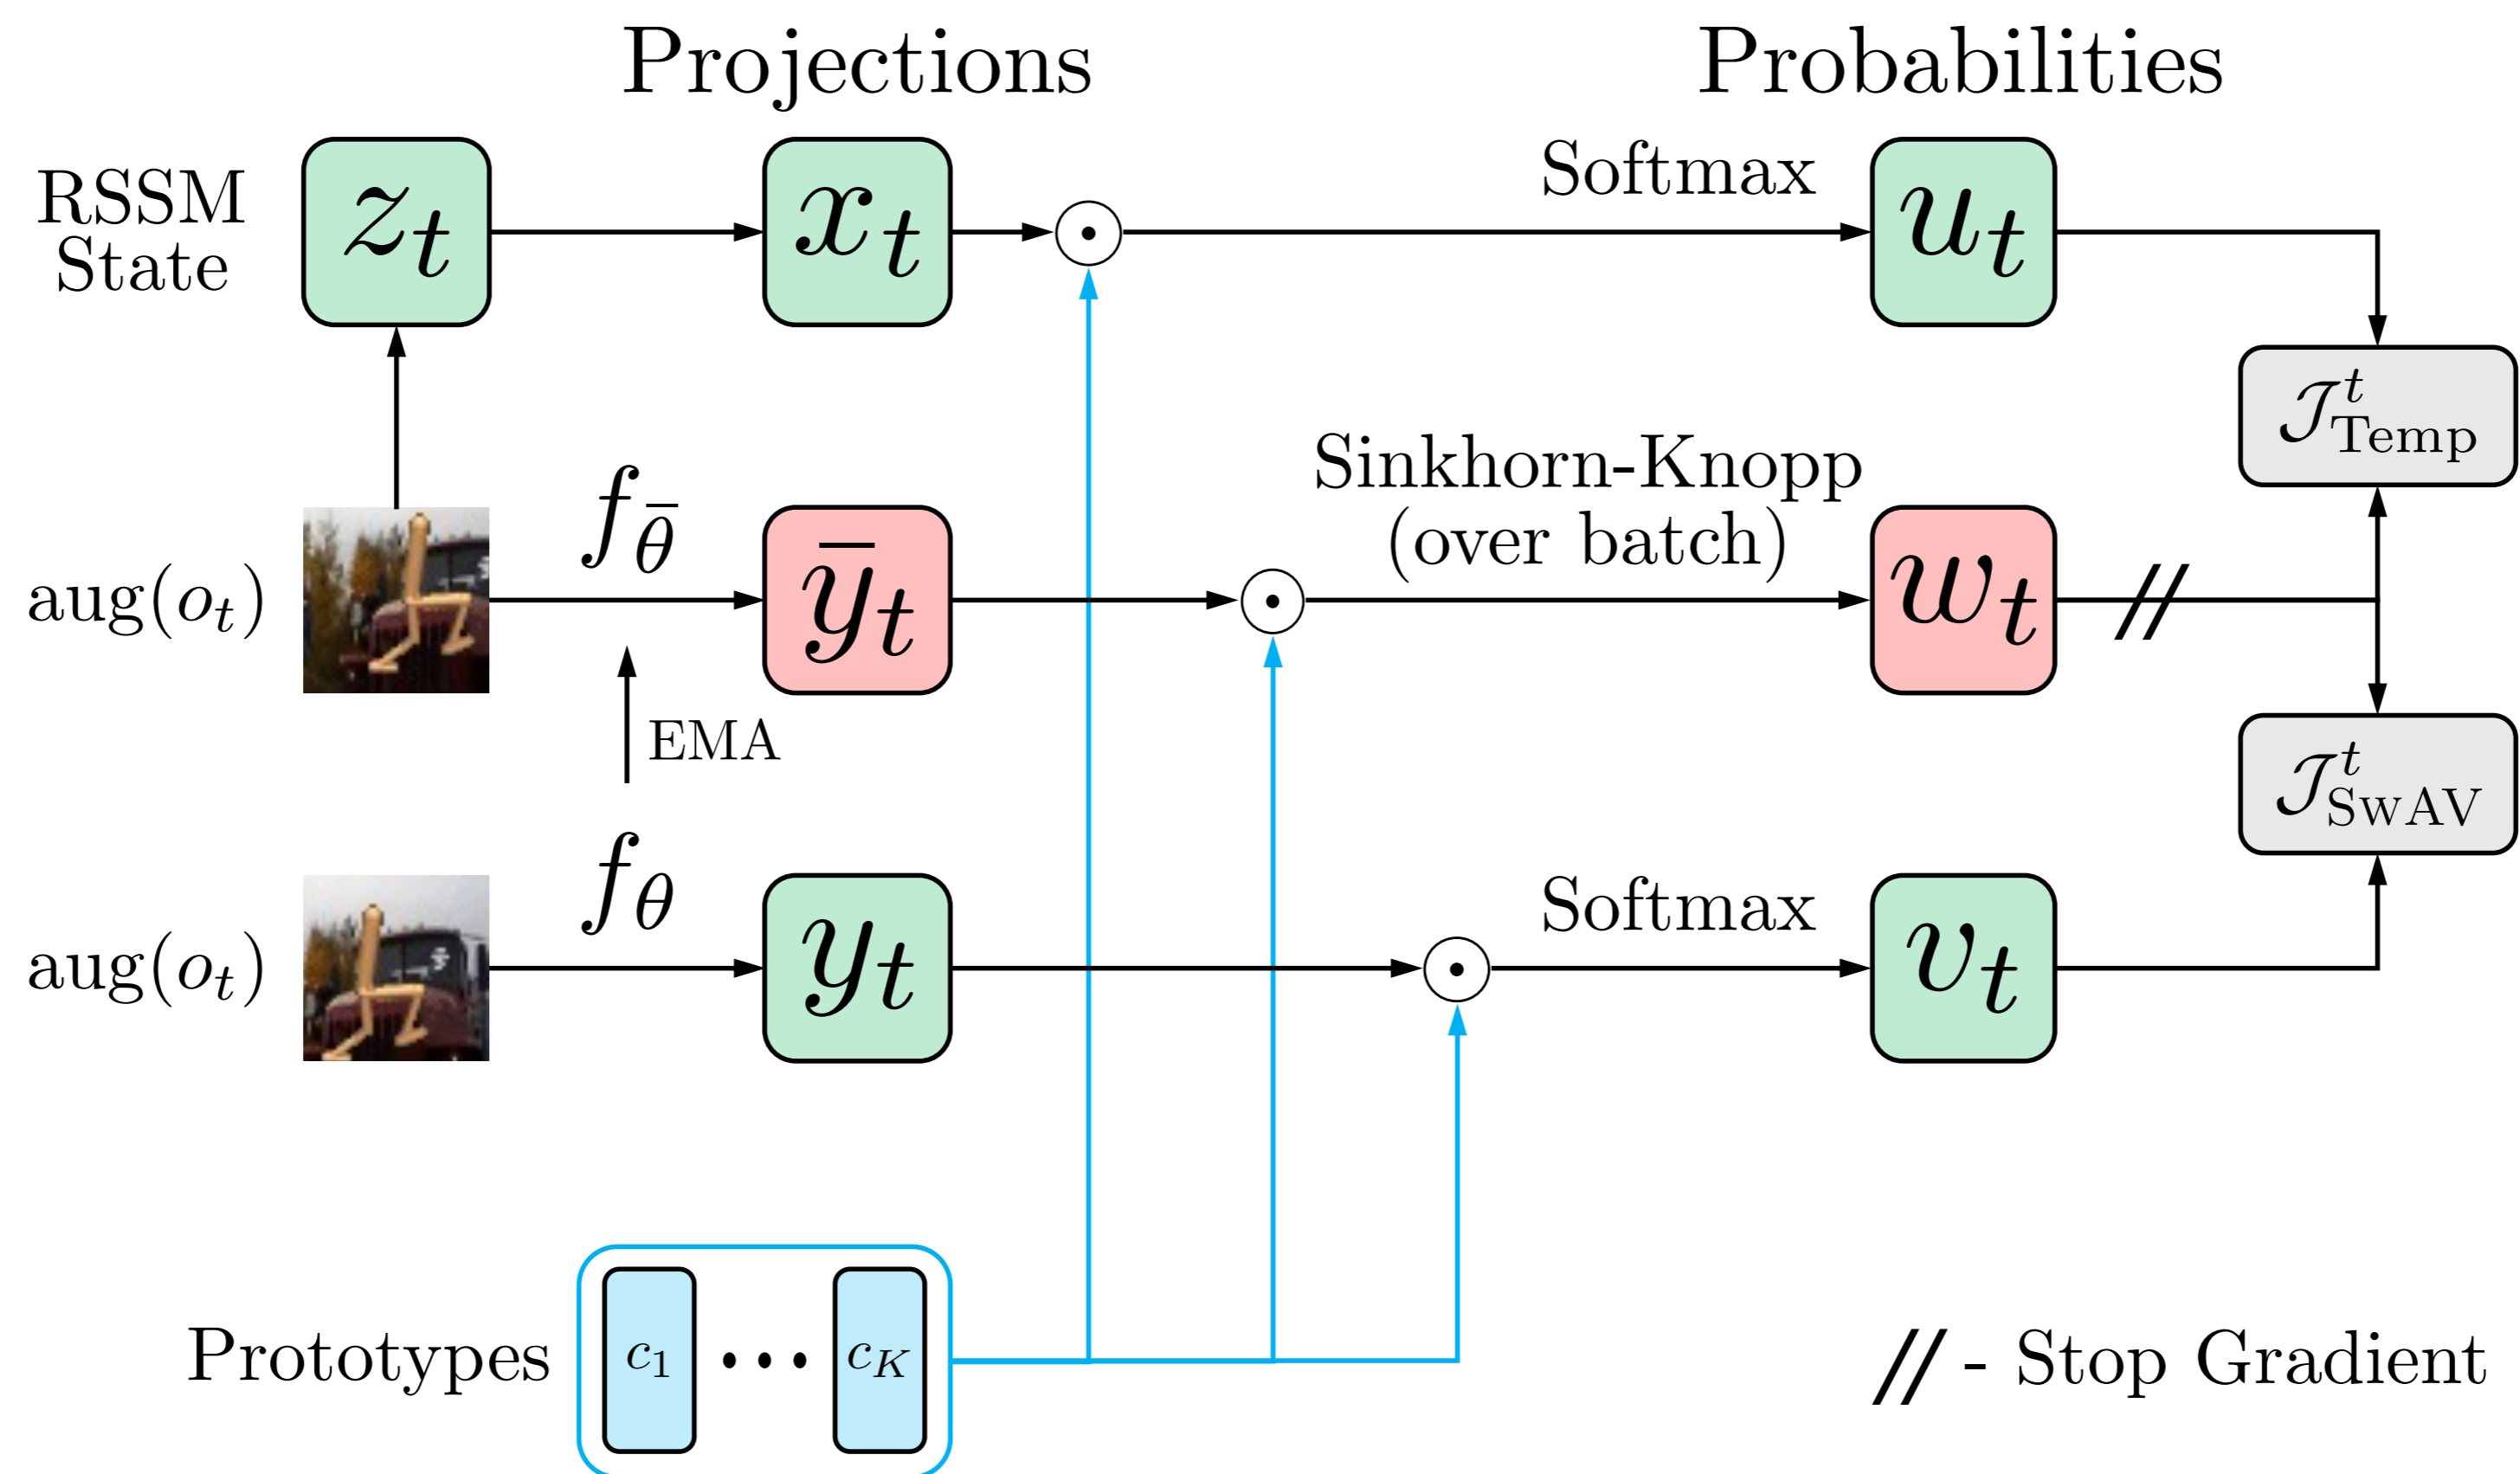


[6] Dream to Control: Learning Behaviors by Latent Imagination. Danijar Hafner et al., ICLR 2020.

[7] Mastering Atari with Discrete World Models. Danijar Hafner et al., ICLR 2021.

[8] Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. Mathilde Caron et al., NeurIPS 2020.

DreamerPro: Dreamer^[6,7] + Prototypes^[8]




[6] Dream to Control: Learning Behaviors by Latent Imagination. Danijar Hafner et al., ICLR 2020.

[7] Mastering Atari with Discrete World Models. Danijar Hafner et al., ICLR 2021.

[8] Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. Mathilde Caron et al., NeurIPS 2020.

Environments



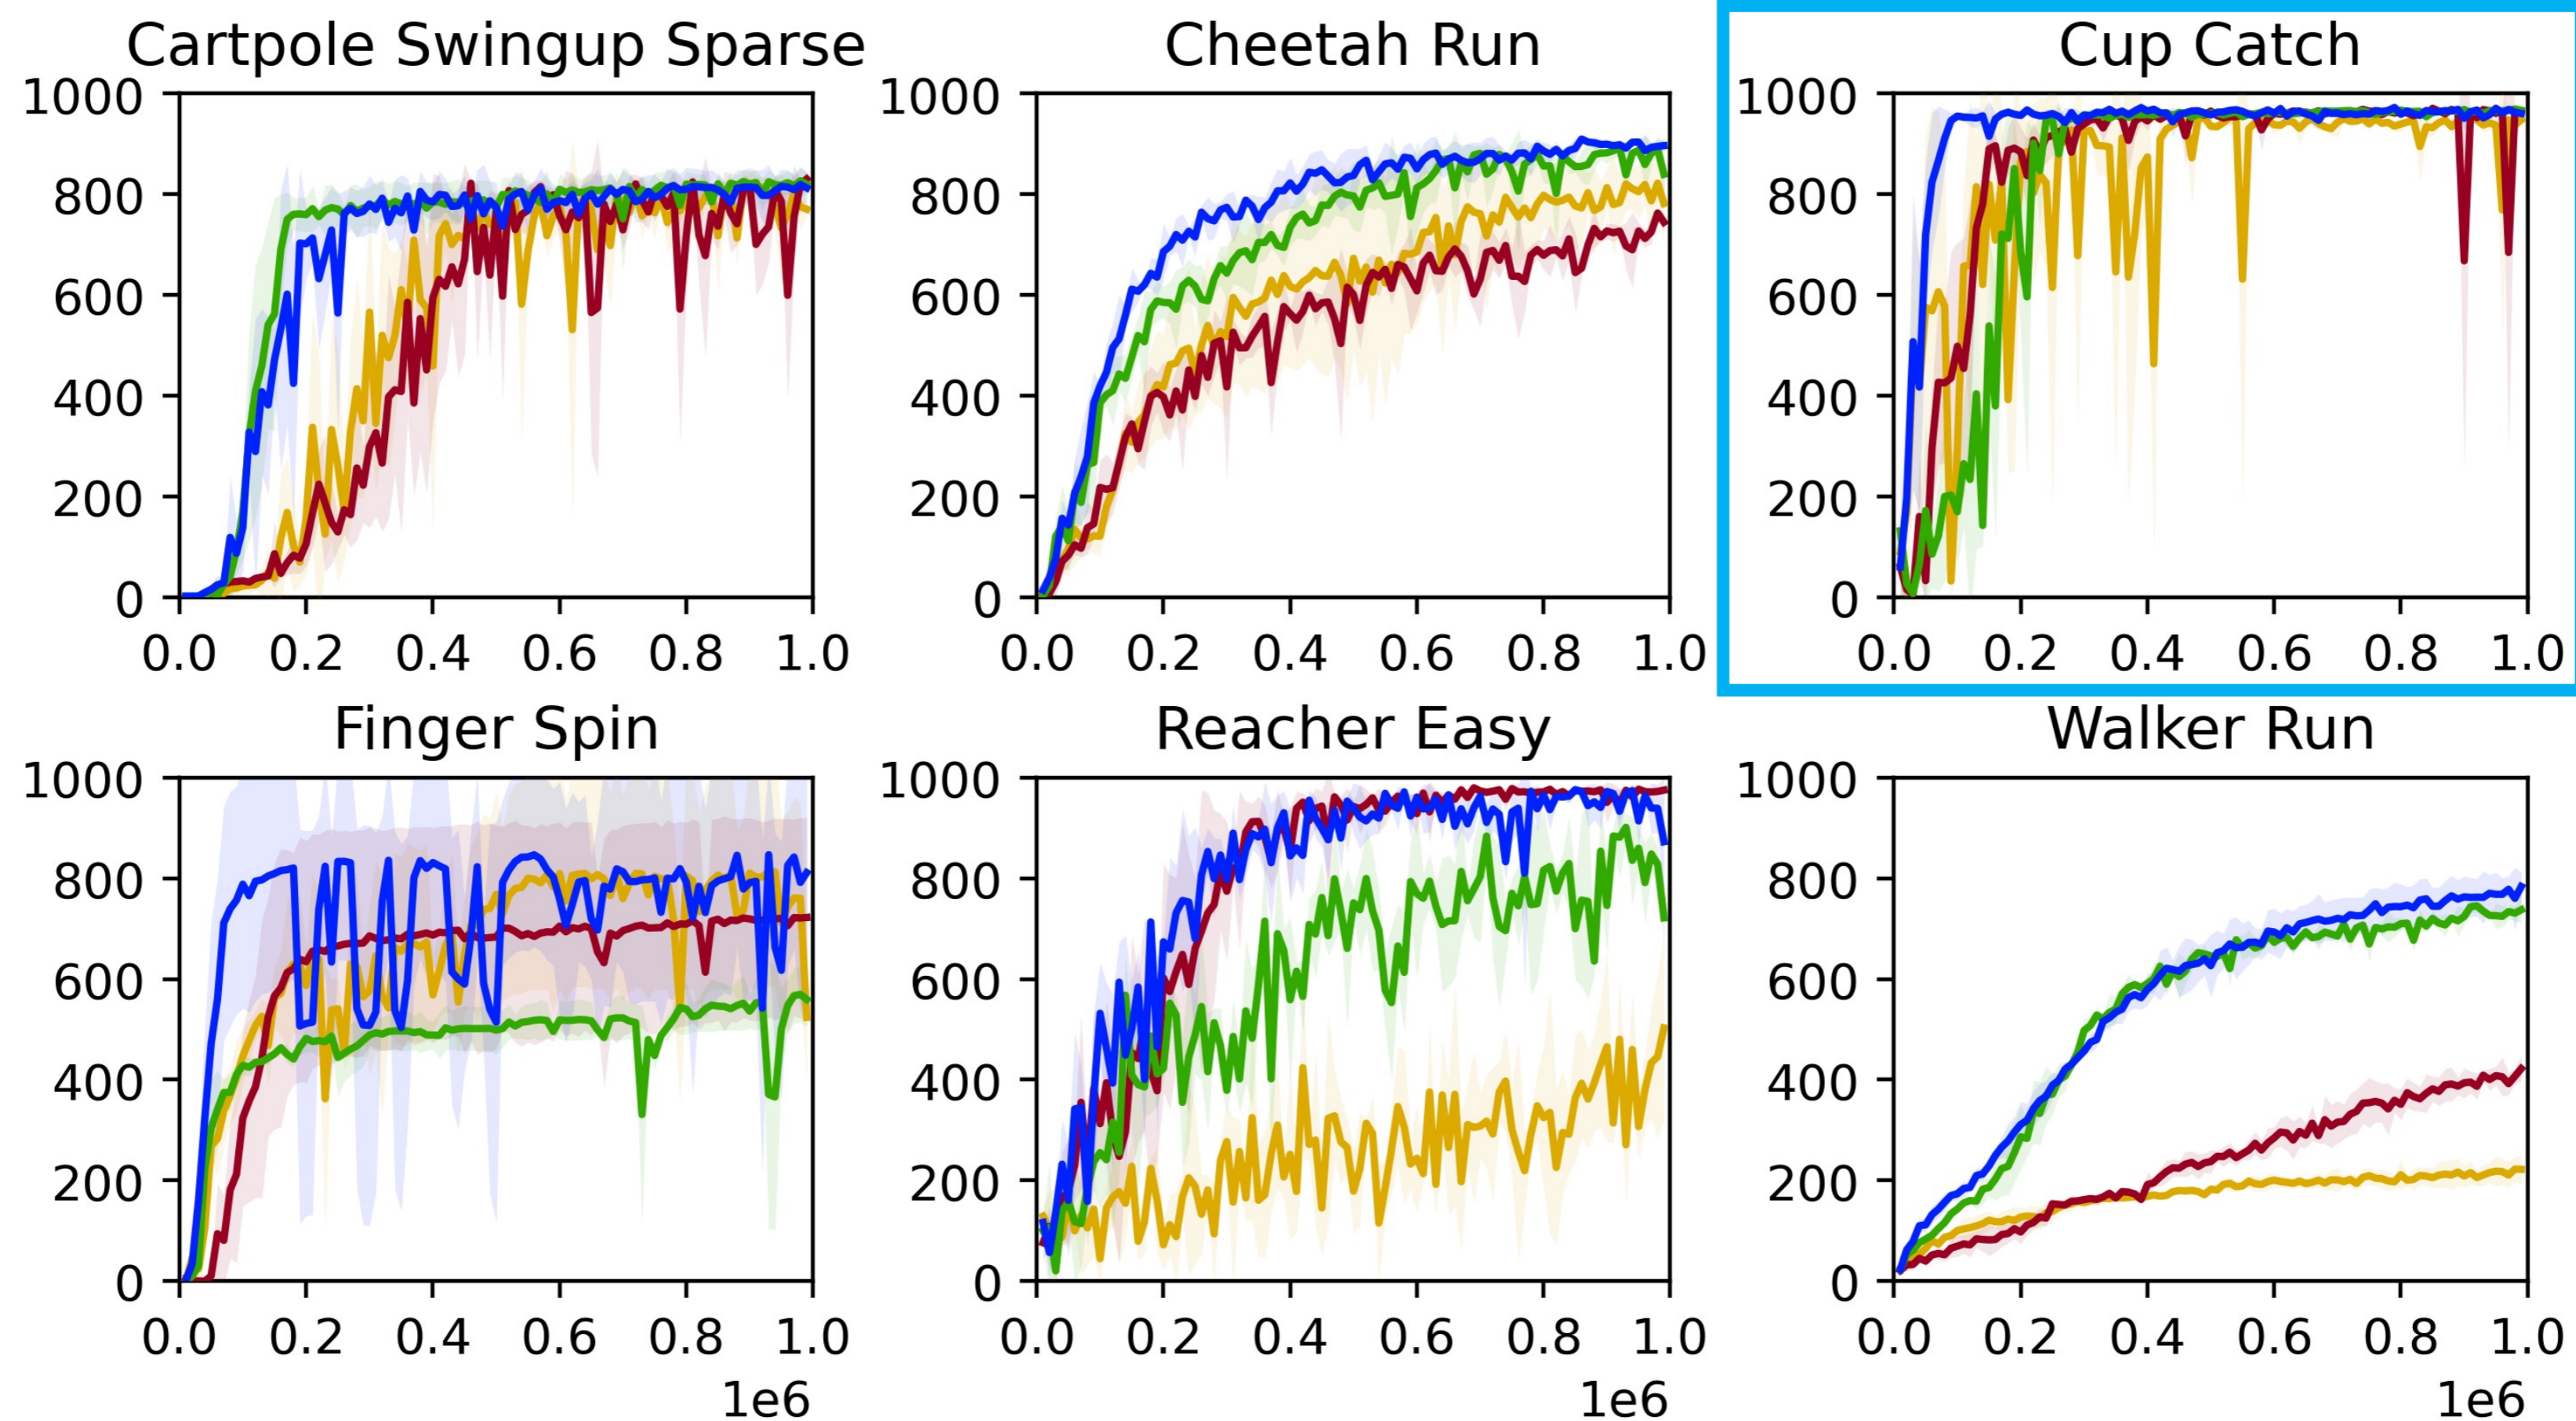
Standard DMC



Natural Background DMC

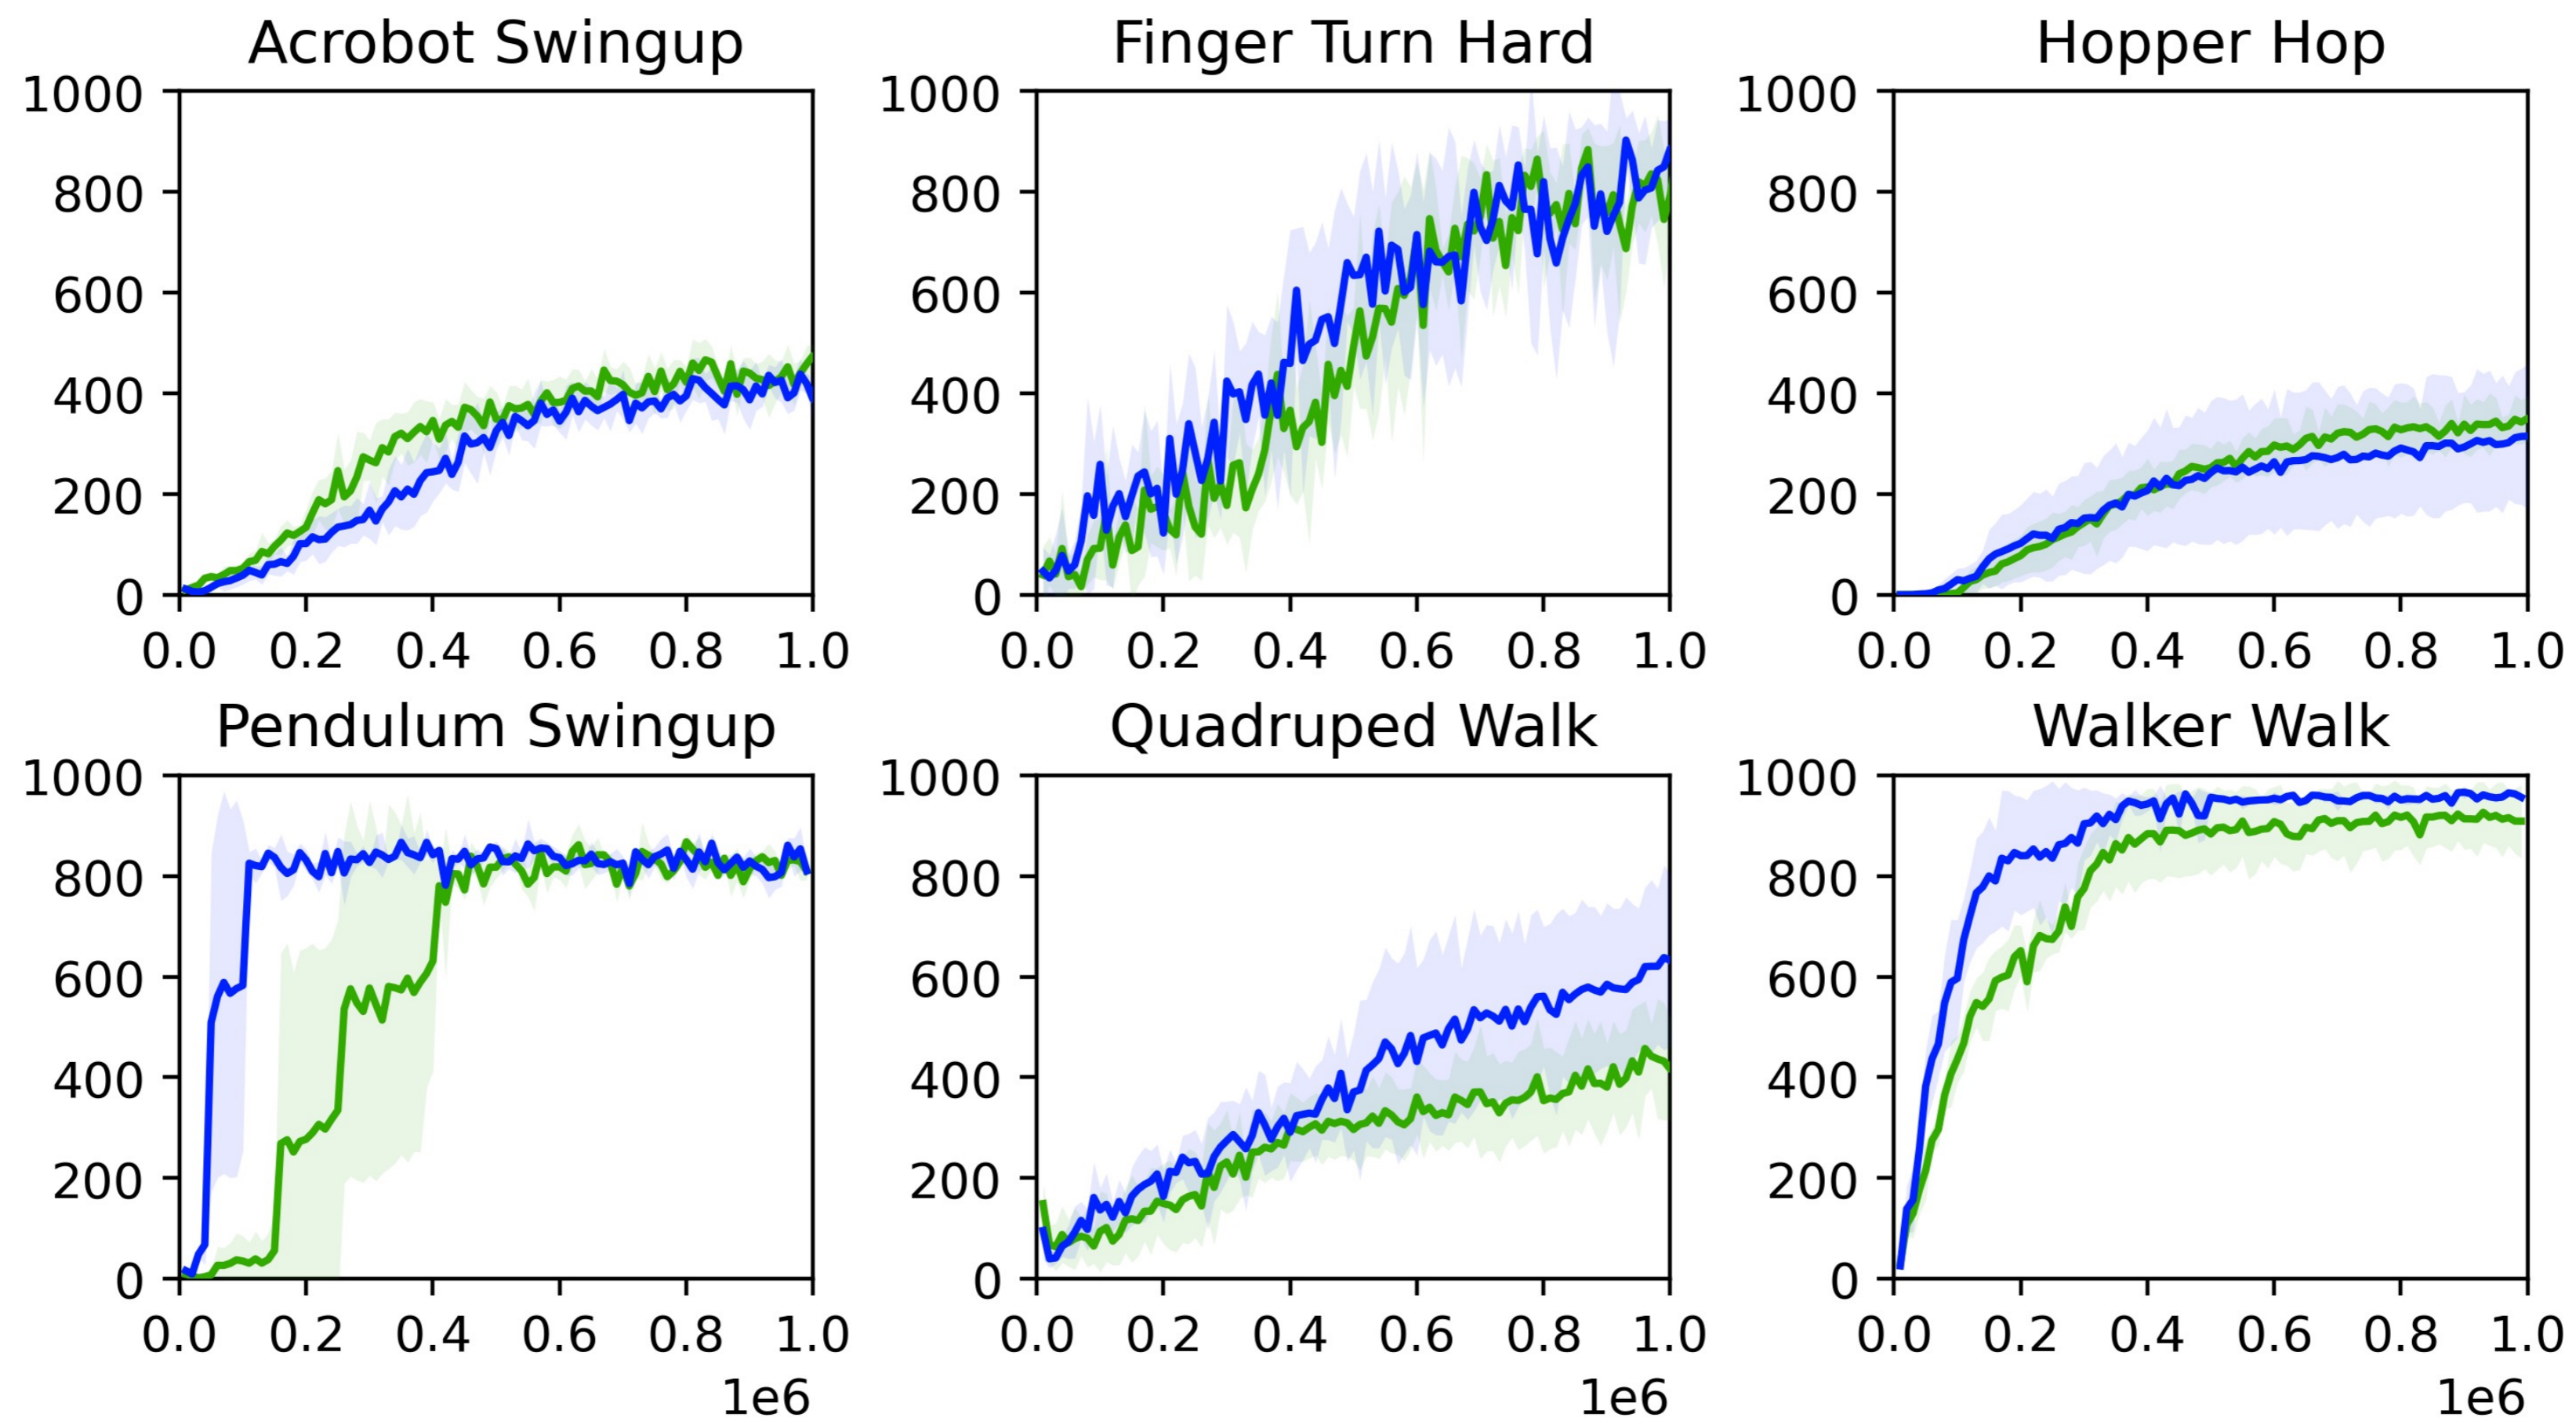
Standard DMC

better data efficiency



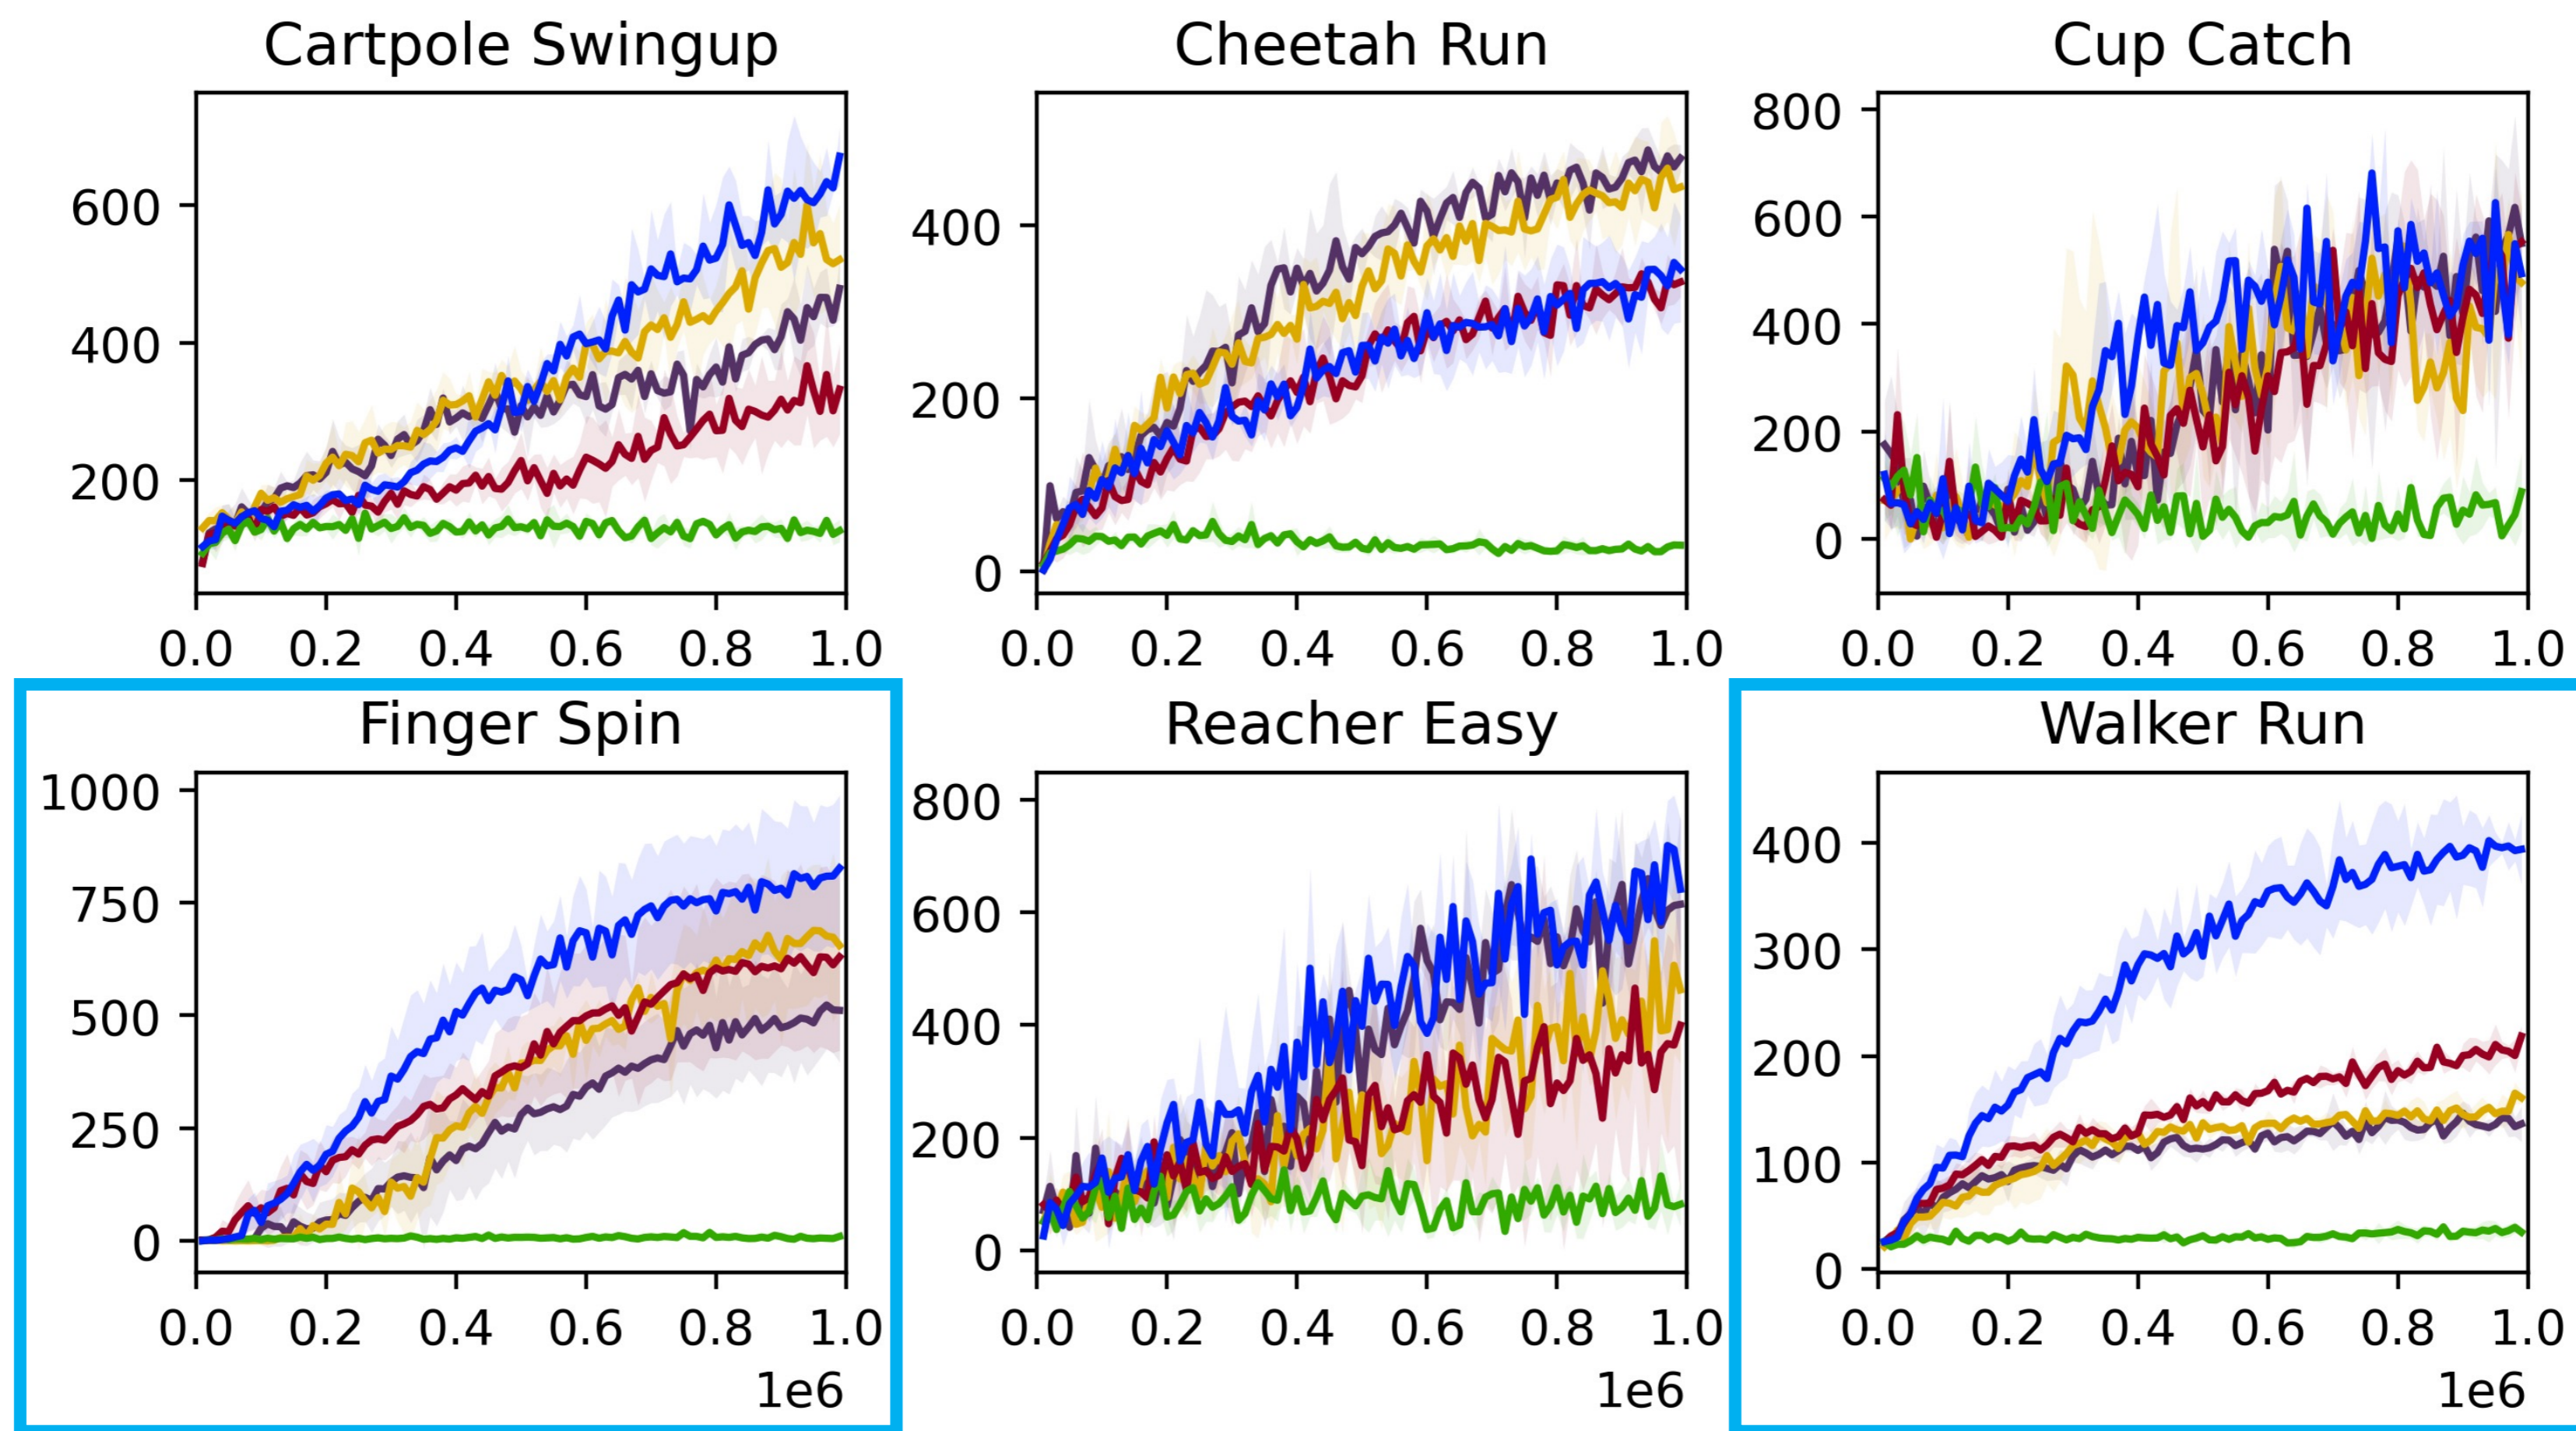
— DreamerPro — Dreamer — Dreaming — TPC

Standard DMC



— DreamerPro — Dreamer

Natural Background DMC



— DreamerPro — Dreamer — Dreaming — TPC — TPC-Batch-300 **large performance gain**

Latent State Visualization




Cup Catch @100K

Input

DREAMERPRO


DREAMER

TPC



Walker Run @1M

Ablation Study



Conclusion

- DreamerPro: First non-contrastive reconstruction-free MBRL agent
- Improve robustness to background distractions without sacrificing performance in distraction-free settings
- Code available at github.com/fdeng18/dreamer-pro